



Performance of the Fast Atmospheric Self Coherent Camera at the NEW-EARTH Lab and a Simplified Measurement Algorithm

Item Type	Conference Paper; Presentation
Authors	Thompson, William; Marois, Christian; Singh, Garima; Lardiere, Olivier; Gerard, Benjamin; Fu, Qiang; Heidrich, Wolfgang
Citation	Thompson, W. R., Marois, C., Singh, G., Lardière, O., Gerard, B. L., Fu, Q., & Heidrich, W. (2022). Performance of the FAST self coherent camera at the NEW-EARTH lab and a simplified SCC measurement algorithm. Adaptive Optics Systems VIII. https://doi.org/10.1117/12.2630415
Eprint version	Publisher's Version/PDF
DOI	10.1117/12.2630415
Publisher	SPIE
Rights	Archived with thanks to SPIE.
Download date	27/09/2023 18:31:29
Link to Item	http://hdl.handle.net/10754/685142

Performance of the Fast Atmospheric Self Coherent Camera at the NEW-EARTH Lab and a Simplified Measurement Algorithm

William Thompson^a, Christian Marois^{b,a}, Garima Singh^b, Olivier Lardière^b, Benjamin Gerard^c, Qiang Fu^d, and Wolfgang Heidrich^d

^aUniversity of Victoria, address, Victoria, Canada

^bNRC-Herzberg Astronomy & Astrophysics, address, Victoria, Canada

^cUniv. of California, Santa Cruz, address, Santa Cruz, United States

^dKing Abdullah Univ. of Science and Technology, address, Thuwal, Saudi Arabia

ABSTRACT

In order to detect low mass and mature planets inwards of approximately 5 AU, future direct imaging instruments will require precision wavefront control that operates at relatively high speed. The self-coherent camera (SCC) is a promising technique for measuring the wavefront from science images at the focal plane. We present here results from NRC's NEW-EARTH lab testing of the Fast Atmospheric SCC Technique, a variant of the SCC and its integration with a Lyot-stop Low-Order Wavefront Sensor. We demonstrate correction of quasi-static speckles in a half dark hole reaching raw 1σ contrasts on the order of 5×10^{-7} at $10 \lambda/D$. We also present a simplified process for extracting measurements and/or DM commands from SCC images using a single matrix-vector multiply. This testing and development are important steps on the way to the upcoming Subaru Pathfinder Instrument for Detection of Exoplanets and Removal of Speckles and the Gemini Planet Imager's CAL2 upgrade.

Keywords: High contrast imaging, focal plane wavefront sensing, self coherent camera

1. INTRODUCTION

In order to directly image faint planets as they orbit their host stars, telescope instruments must correct the incoming wavefront with high accuracy. The two primary sources of wavefront error are atmospheric speckles, caused by our turbulent atmosphere's changing index of refraction, and instrumental, also known as quasi-static speckles.¹ Once the wavefront is corrected, a focal plane mask (FPM) and Lyot stop are typically used to block the light from the star while letting through as much light from the planet as possible.

The Self Coherent Camera, or SCC, is a technique for focal plane wavefront sensing in the context of coronagraphy pioneered by Baudoz et al.² Rather than merely discarding diffracted light at the Lyot stop, a small amount of the light is allowed to pass through a second much smaller aperture placed off-axis, referred to as the "pinhole". This light creates a reference beam that propagates to the focal plane where it interferes much like a traditional double-slit experiment with any starlight that makes it past the FPM. Given that the pinhole is quite small, it acts as a filter that produces a near-perfect reference beam that only passes a small amount of low-order aberration.

The resulting fringe patterns between the main beam and the reference beam encode the relative phase offset between the two arms of the common-path interferometer. This information can be processed to extract the full complex electric field at the focal plane to enable active control and post-processing with Coherent Differential Imaging (CDI³⁻¹⁰).

Further author information: (Send correspondence to W.T.)

W.T.: E-mail: wthompson@uvic.ca

C.M.: E-mail: christian.marois@nrc-cnrc.gc.ca

Since the light in the reference beam picked off by the FPM comes only from the star, the resulting interference pattern is only present between the reference beam and residual starlight. Light from a planet or disk travels through only one arm of the interferometer (the main pupil) so is not fringed.

A key improvement to this technique was developed by Gerard et al.^{9,11} the Fast Atmospheric Self-coherent-camera Technique (FAST). In FAST, the FPM is modified to direct the stellar core further off-axis and concentrate the light directly on the pinhole. This increases the light through the pinhole significantly to allow a better match between the intensities of residual speckles and the reference beam. The result is that fringe pattern can be detected in the same amount of time as the speckle itself, dramatically increasing the possible correction rate.

In this work, we present results of our laboratory testing of FAST: both active DM correction and CDI post-processing. We also present an algorithm for extracting SCC measurements that is numerically faster, tolerant of saturated pixels, and simpler to implement. This work is an important step towards the Subaru Pathfinder Instrument for Detection of Exoplanets and Removal of Speckles (SPIDERS)^{12,13} and the Gemini Planet Imager's CAL2 upgrade.¹⁴ Both of these instruments under development will include an SCC with very similar configurations to the laboratory setup we present here. Achieving not only the best possible performance from the SCC, but also robust and reliable operation are critical to meeting the science goals of these projects.

2. EXPERIMENTAL SETUP

We have been working to test FAST at the NEW-EARTH (NRC Extreme Wavefront control for Exoplanet Adaptive optics Research Topics at Herzberg) laboratory. The following describes our experimental setup in brief, and the full details are available in a previous conference proceeding by Lardière et al.¹⁵

The VIPER (Victoria Pathfinder for Exoplanet Research) bench in the NEW-EARTH lab contains: an ALPAO 468 deformable mirror, tilt-gaussian-vortex (TGV) FAST FPM, reflective Lyot stop with pinhole for FAST, Lyot-stop Low-Order Wavefront Sensor (LLOWFS^{16,17}), and CRED 2 science camera / focal-plane wavefront sensor. The TGV FPM is a fresnel design manufactured using additive lithographic fabrication by KAUST with four discrete layers. The bench also includes an imaging Fourier transform spectrograph to be described in future papers. The tests presented in this proceeding used an unpolarized white SuperK COMPACT light source from NKT Photonics with a 2% bandpass filter rather than a narrow laser. This reproduces the slight blurring of fringes in a realistic bandpass.

The LLOWFS arm is responsible for controlling low-order zernike modes and keeping the star centred on the FPM. It sees all photons that are blocked by the coronagraph, so has access to significantly more light and accordingly can run at similar speeds to traditional Adaptive Optics loops. The CRED-2 is placed downstream of the Lyot stop where it functions as both the science detector and focal-plane wavefront sensor. The FPM is designed such that the reference pinhole is quite far off-axis. The resulting fringe pattern between main pupil and the reference beam is at high spatial (less than $1/4 \lambda/D$). This fully separates information in the fringes from any astrophysical signal.

The ALPAO-468 Deformable Mirror (DM) has a circular arrangement of actuators providing us with a $24 \lambda/D$ diameter control region. In its current setup, the reflective Lyot stop blocks the two outer annuli of actuators. These actuators can therefore not be sensed by the SCC, but are still seen with limited sensitivity by the LLOWFS (proceeding of Singh et al 2022¹⁸). We operate the SCC in both full and half dark hole modes, allowing us to correct phase errors across the entire control region or phase and amplitude errors across any one half of the control region.

We placed a Thorlabs optical chopper behind the reflective Lyot stop's reference pinhole. This allows us to record fringed and unfringed pairs of images. With the pinhole placed off-axis, it is possible to operate the SCC without a chopper; however, we found that operating on differential images improves the stability of system and allows us to safely increase the gain. The current optical chopper is limited to at lowest a 200Hz on/off chopping rate. The SPIDERS pathfinder described in proceedings of Lardière et al.¹² and Marois et al.¹³ will employ a custom-built optical chopper that can smoothly operate at lower speeds (see Johnson et al. in these proceedings¹⁹).

3. SCC: MEASUREMENT ALGORITHM AND RESULTS

We now describe how we implemented the SCC control loop on the VIPER bench and a simplified algorithm used for SCC measurement.

3.1 Calibration

Just as in previous SCC implementations, we calibrate the SCC using Fourier modes applied to the DM. It is possible to calibrate an SCC with any combination of DM commands as long as they form a complete and independent basis. In practice, Fourier modes are one of the best choices of calibration basis since they concentrate the impact of the command in just two $\sim \lambda/D$ spots. This increases the SNR of the poke and speeds up calibration compared to, for example, poking individual actuators. Since the outer two annuli of actuators are blocked by the Lyot stop, we do not control those actuators with the SCC. We also remove any tip, tilt, and piston components from the Fourier modes.

We choose the amplitude of the calibration Fourier modes to be as small as possible without making the calibration sequence impractically slow. Smaller amplitude calibrations should improve the linearity of the SCC²⁰ but also decrease the calibration signal and thus increase the integration time needed during calibration. We apply each Fourier mode in a push-pull sequence. Figure 1 shows the fringed and unfringed images recorded for a push and pull of one mode as well as several combinations of these images. It is sufficient to use only fringed images in the calibration (i.e., push - pull, as outlined by the red box) however, we find that in practice the DM does not apply a perfectly symmetric change in the wavefront between push and pull. This results in some unfringed structure which is undesirable in our calibration matrix. Instead, we calibrate the SCC by calculating

$$(\text{push, fringed} - \text{push, unfringed}) - (\text{pull, fringed} - \text{pull, unfringed}) = \quad (1)$$

$$\text{push, fringed} - \text{pull, fringed} - \text{push, unfringed} + \text{pull, unfringed} = \quad (2)$$

$$(\text{push, fringed} - \text{pull, fringed}) - (\text{push, unfringed} - \text{pull, unfringed}). \quad (3)$$

This does a better job isolating just the change in fringes (Figure 1, outlined by blue box). The remaining non-fringe structure is only a result of turbulence or changing source intensity, and is typically only visible near the core. We finally suppress this structure using a high-pass filter of λ/D (or equivalently, fourier filtering to mask out the center of the modulation transfer function). Note however that this filtering is only performed during calibration, and would also not be required in a perfectly stable environment.

In order to show the Fourier modes clearly in Fig. 1, we first applied a static full dark hole correction. This isn't necessary in general, though calibrating iteratively in this way may improve stability when bringing up a system with significant non-common path errors. A downside of this calibration and measurement algorithm is that it does not separate out and track the reference beam PSF. A changing reference beam PSF would have the effect of spatially modulating the gain of the SCC. For systems where the reference PSF is not stable, it could still be estimated using differential images with Fourier filtering or fit using a PSF model. Here, we assume that any reference PSF during on-sky operation is on average the same as the one used to capture the calibration sequence modulo changing source intensity. This assumption should hold on-sky thanks to the LLOWFS.

Figure 2 shows the results of this calibration procedure for three Fourier modes. The sub-figure **B** corresponds to the images shown in the previous figure. In previous works²⁰ it was recommended to apply a circular mask around bright peaks to reduce crosstalk between modes. We do not currently apply those masks to our calibration matrix, but future testing with atmospheric residuals is warranted to fully evaluate their impacts.

In order to select the SCC's correction region, we create a mask with the same dimensions as the image with either a full dark hole or a chosen half-dark hole. With the TGV FPM, we limit this mask to begin at $2\lambda/D$ in separation from the star and end at approximately $11\lambda/D$. It is important to fully exclude pixels outside the control region because the DM, while not having the necessary actuators to properly correct areas outside the control region, can still apply small corrections using second-order effects (e.g. second order sine waves out to $24\lambda/D$ visible in some calibration images). Rather than finding a compromise between a very poor correction outside the control region and a slightly worse correction inside the control region, we prefer that the SCC focus entirely on the control region where a full correction is possible.

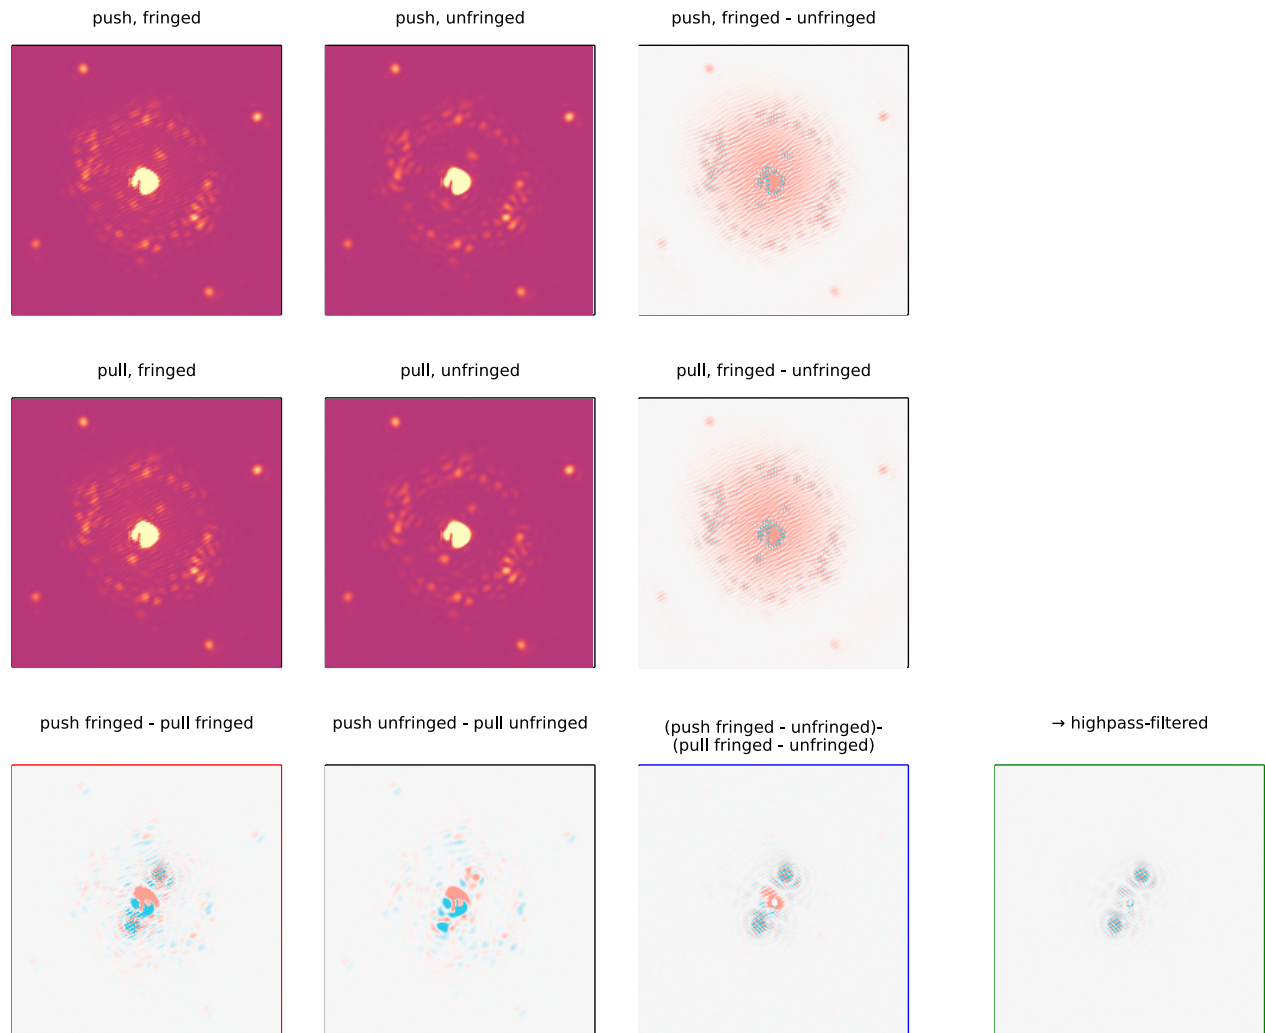


Figure 1. Images illustrating the calibration of FAST at the NEW-EARTH lab’s VIPER bench. In these images, a slice of the interaction matrix for one Fourier mode is recorded (outlined with a green box). The push and pull differential images (third column, top and middle) also show the envelope of the reference PSF (outlined with a red box). The four spots visible outside the control region are caused by print-through on the DM membrane.

To generate the SCC interaction matrix, we flatten each image and select just the masked pixels in our desired dark hole region. We then concatenate them into a $N_{\text{modes}} \times N_{\text{selected pixels}}$ matrix and take the pseudo-inverse. In order to regularize the matrix inversion and reduce noise propagation, we employ a Tikhonov regularization. Finally, this inverted matrix is multiplied with a matrix consisting of the commands used to calibrate the SCC. This results in a $N_{\text{pixel}} \times N_{\text{actuator}}$ matrix.

As an additional step to prevent noise propagation, we calculate the sensitivity of SCC towards each mode. We then apply a cutoff to completely remove modes that are poorly sensed by the SCC— typically those that fall at the edges of the control region or SCC mask (to be described in the following section).

3.2 Measurement

Once we have the SCC’s interaction matrix, we can take a measurement from the SCC using a single matrix-vector multiply. We work with differential images (fringed - unfringed). Operating on differential images removes the effects of changing detector cosmetics and dark current. As an additional benefit, the unfringed frames can




Figure 2. Prepared slices from the interaction matrix for three different sinusoidal modes. The DM is not capable of perfectly reproducing sinewaves so there is some noticeable cross talk (central region in **C**), particularly with sinewaves much greater than 6 cycles/pupil. We do not apodize the beam in VIPER so the Fourier modes show Airy rings.

serve as science images since they are devoid of photon noise from the reference beam. In our experience the reduction in correction rate is a worthwhile trade-off to increase loop stability. As an additional benefit, we can choose after the fact to use both the fringed and unfringed frames as our science images if we are speckle limited or only the unfringed frames if we are limited by photon noise of the reference beam.

First, we collect into a vector the pixels in the masked dark hole region (as was chosen when generating the interaction matrix). Next, we multiply this vector of pixels by the interaction matrix described above to get a vector of DM actuator commands. Finally, we employ a basic integrator in closed loop. The DM command is updated by:

$$\text{cmd} = \text{cmd} - g \times \text{measurement}$$

We apply this algorithm to differential images but it is perfectly valid to use it on fringed images only. This is because the fringes isolated by the above calibration procedure are symmetric in intensity around zero: any smooth structure is removed by the push/pull fringed/unfringed double difference and by the high-pass filter. As such, they can only fit the fringes in an image and not any smooth incoherent structure.

This algorithm is computationally efficient, consisting only of a pixel gathering or masking operation followed by a single matrix-vector multiply. Additionally, by avoiding Fourier transforms, it is robust to saturated pixels and some detector artifacts. In differential images, saturated regions become zero and simply contribute no information to the SCC. If applied to fringed images instead of differential images, the same holds because saturated regions cannot be fit using a fringe pattern with symmetric intensity.

Robustness to saturated pixels is important for a few reasons. First, all real detectors have some level of read noise so it is necessary to increase the detector's exposure time until the fringe signal of a typical speckle is above this noise level. At the same time, physical detectors, especially those with very low read noise, have limited dynamic range. For example, a 14-bit analogue to digital converter (ADC) has roughly 4 orders of magnitude dynamic range. If for algorithmic reasons we were prevented from ever saturating the detector, then considering a FPM that suppresses core light by least $100\times$, the faintest easily detected speckle would be just 10^{-6} . In our particular implementation, we limited spatial frequency of modes controlled by the SCC. Approaching the edge of the control region, the sampling of the DM falls to just under one sinewave every two actuators. This makes it challenging for the DM to accurately construct the sine waves at higher spatial frequencies. We found that by limiting the DM's correction to an outer separation of $10\lambda/D$ (or even 6) results in improved contrast closer in.

Since the slowest changing quasi-static speckles have the largest impact on direct imaging performance in most conditions, we aim for the SCC to provide a strong DC correction. As such, we eschewed using a leaky integrator. Without a leak, this implementation was more susceptible to poorly sensed modes building up in the

integrator. To prevent this kind of instability, we were careful to remove modes that were poorly sensed by the SCC.

One unexpected noise source was evolving dark current patterns on the CRED-2 detector. The detector exhibits a pattern of columns that slowly change in intensity over a period of seconds to minutes. This makes them difficult to remove with a dark calibration. Working with differential images does remove most of this evolving dark current but still leads to a small net linear drift recorded by the SCC. If this signal propagates into a mode that is poorly sensed by the SCC, a command can slowly build up in the integrator.

To limit their effect, we rotated the detector by 30 degrees. This ensures the column pattern (and a much fainter row-oriented pattern) do not match the orientation of the fringes. As a further mitigation to evolving detector bias/dark current, we intend in future setups to alternate the fringed - unfringed subtraction sequence in time. This will use each frame twice at twice the speed and the double-difference will remove any net linear drift in the detector bias.

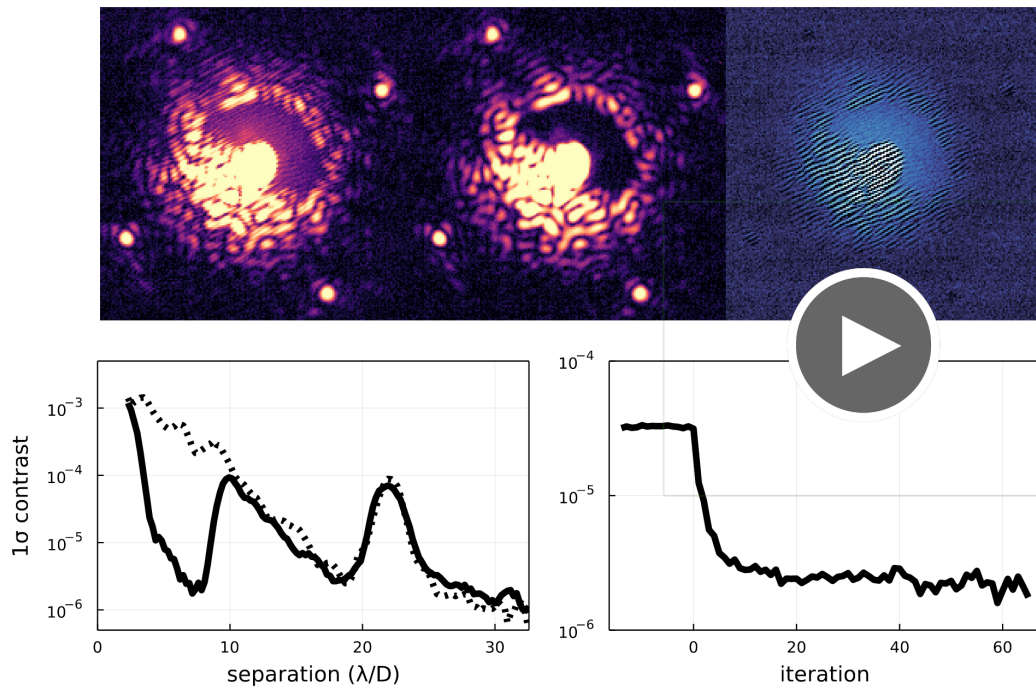


Figure 3. Video 1. A video of the SCC loop closing on the bench's static aberrations and local bench turbulence. **Left:** fringed SCC image. **Center:** unfringed SCC image. **Right:** differential SCC image. **Bottom left:** contrast in the half dark-hole side of the unfringed image as a function of separation at the current iteration. **Bottom right** best contrast in the half dark hole of the unfringed image over time. Some faint incoherent structure is visible above the star in the fringed and unfringed image. These individual short exposure images are limited by read noise so the contrast appears to plateau after around 15 iterations. <http://dx.doi.org/10.1117/12.2630415.1>

3.3 Laboratory Results

We achieved promising results with the SCC on the VIPER bench. Video 3 shows an animation of closing the SCC loop on static errors at a gain of 0.3. This video was slowed down by a factor of 10 so that individual iterations are visible. The loop corrects the half dark hole down to nearly the level of read noise in less than 10 iterations (with the bulk of the correction taking just 2 iterations). In this video, the contrast is measured per frame at 400 frames per second. Accordingly, the contrast inside the half dark hole plateaus near the limit of the detector read noise.

Figure 4 shows the contrast on the bright and dark halves of the image for a stack of 300 unfringed frames for a total of 300 ms integration. In this time, we achieve a contrast of $\sim 1 \times 10^{-6}$ inside the dark hole before CDI, or roughly a $500\times$ improvement compared to the light side.




Figure 4. Stack of unfringed SCC images with a higher source intensity than Figure 3 shown with an arcsin stretch.

Finally, we show in 5 how the SCC loop generates a full dark hole despite starting with most of the control region saturated thanks to the improved measurement algorithm. The SCC does not have access to any information from saturated pixels so the correction must proceed inwards from unsaturated areas. This process is slow (it takes on the order of 100 iterations) but robust and, unlike algorithms that require Fourier transforms, the loop remains stable.

4. POST PROCESSING WITH COHERENT DIFFERENTIAL IMAGING

In addition to real time correction, the SCC allows us to reconstruct the stellar PSF at each iteration. This enables a type of differential imaging in the spirit of reference star differential imaging (RDI²¹), angular differential imaging (ADI²²), and spectral differential imaging (SDI^{23–25}) called coherent differential imaging (CDI). CDI is possible because the SCC fringes allow us to differentiate between light that travelled through only the main pupil (unfringed in the science images) and light that travelled through both main pupil and the reference pinhole (fringed in the science images).

CDI as a current concept was developed over a series of papers^{3–10} that are well-summarized in the thesis of B. Gerard.²⁰ Simulations from some of those works indicated that CDI using FAST could approach roughly $100\times$ improvement on a 0th magnitude star. In CDI, one reconstructs the stellar PSF using wavefront measurements from the SCC. The SCC only measures coherent light from the star so this reconstruction is free of contamination from other astrophysical sources. Since these wavefront measurements are done on the science detector it is in principle easy to match the reconstructed PSF with the target image.

One unexpected challenge we have encountered in achieving these simulated CDI results with the SCC is the interpretation of the sidelobe as a shifted copy of the stellar electric field from the main pupil. In fact, the sidelobe and main pupil differ because they are measured by the detector on different spatial scales, and thus are convolved by different parts of the pixel modulation transfer function (MTF). The pixel MTF is a broad sinc function that results from the finite size of pixels in the focal plane. Effectively, detectors are less sensitive to very high spatial frequency structures than low spatial frequency structures. Effectively, we must reconstruct not just the wavefront but the wavefront as sampled by the science detector. Since we are aiming to subtract speckles with a high degree of accuracy, this complication from the pixel MTF quickly becomes significant.




Figure 5. Sequence in time (iteration increasing downwards) of the SCC converging on a full dark hole despite having more than half the control region initially saturated. From left to right: fringed, unfringed, and difference. Convergence is slow in saturated regions since the SCC must work its way inwards from surrounding unsaturated pixels, but the loop nonetheless makes progress and remains stable.

4.1 Fourier Transform and Model Based Algorithm

The first solution we present is to begin by deconvolving the science images by the pixel PSF. This leads to a slight sharpening of the science image and a frequency dependent amplitude correction of the sidelobe used to construct the reference image. For CRED2, we modelled the pixel MTF as a sinc function (Figure 6). We captured a sequence of fringed and unfringed images and optimized the radial extent of the pixel MTF to get the best subtraction. As expected, the best subtraction occurred when the sinc function was scaled to have its first null at the far corners of the MTF. Instead of adopting this model based approach, future work could directly measure the pixel MTF using a test pattern or knife-edge test. The results of this modeling are shown in Figures 7 and 8. The downside of this approach is that it requires a deconvolution step. Deconvolutions are often avoided in image processing since they can propagate noise across spatial frequencies.

The reference pinhole PSF is required in order to reconstruct the stellar PSF from the differential images. This pinhole is not infinitely small, so the resulting PSF is affected by some low-order aberrations that may evolve throughout a sequence. Since we use an optical chopper, it's possible to extract the pinhole reference




Figure 6. Best fitting pixel MTF and corresponding pixel PSF of our CRED2 on a log scale. The outline of a fringed image MTF is plotted overtop the pixel MTF in white. This shows how the sidelobes are convolved by a different, non-symmetric part of the pixel MTF than the main pupil.




Figure 7. FFT based CDI subtraction. Left: unfringed image. Right: post-processed image.

PSF from the low frequency components of the differential images. However, we find that this signal is quite noisy in individual pairs of images. We instead use a synthetic model of the reference PSF as an Airy disk (with a possible central obscuration) and optimize it's location, scale, and position for sub-sequences of fringed and unfringed images. In our Julia-based software, this is done easily using the package `PSFModels.jl`^{*}.

4.2 Self-Consistent Image Plane Algorithm

The CDI solution described above is effective, but like the original SCC measurement algorithm, it requires a series of Fourier transforms that may propagate noise from saturated pixels. It furthermore requires a deconvolution step which is not ideal for noisy data.

A future approach we aim to test in the NEW-EARTH lab is an image domain algorithm in the same vein as the new SCC measurement algorithm presented in this paper. Our proposed CDI approach is to leverage the SCC calibration matrix inside the control region of the DM. We would first apply the image-based SCC measurement algorithm to find a new modal basis. Using this modal measurement, we would then compute a linear combination of unfringed sine wave spots recorded simultaneously with the SCC calibration matrix.

^{*}<https://github.com/JuliaAstro/PSFModels.jl>




Figure 8. Contrast improvement shown in Figure 7.

That is,

$$((\text{push, unfringed} - \text{neutral, unfringed}) - (\text{pull, unfringed} - \text{neutral, unfringed})) \quad (4)$$

Effectively, we would directly match the fringes in an image with unfringed speckles. Since we would perform this correction in software rather than with a hardware DM, we could add additional flexibility by splitting the two spots from the DM Fourier spots into independent modes. This PSF reconstruction would self-consistently incorporate the reference beam PSF and pixel MTF. Two limitations of this approach would be that it only works within the control region and that it does not respond to a changing reference beam PSF. Thankfully this PSF is spatially filtered and in our case, stabilized by the LLOWFS. In order to expand the region of subtraction outside the control region of the DM, it may be possible to use on-sky sequences where turbulence naturally produces speckle variations at these separations. To build such a CDI speckle library, a second form of diversity would be required like ADI²² or RDI²¹ to ensure that any planets or other astrophysical sources do not contaminate the calibration. This could allow the image-based CDI subtraction algorithm to function over the full field of view illuminated by the reference beam.

Though not yet fully developed, an example of this approach is presented in Figure 9. We aim to further test this approach at the NEW-EARTH lab in the coming months.

5. CONCLUSION

We have presented the current performance of FAST in the NEW-EARTH lab, and in addition, improved algorithms for extracting SCC measurements and CDI reconstructions. On the VIPER bench, we have routinely achieved raw 1σ contrasts of 5×10^{-7} in relatively uncontrolled conditions by combining the SCC and CDI. These improved algorithms will be important for robustly implementing FAST on-sky in the SPIDERS pathfinder and Gemini Planet Imager CAL2 upgrade.

ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), 466479467. This research used the facilities of the Canadian Astronomy Data Centre operated by the National Research Council of Canada with the support of the Canadian Space Agency. The authors are grateful to have performed this work on the traditional territories of the Coast Salish and Lekwungen-speaking peoples of Vancouver Island.




Figure 9. Initial results of an image-domain based CDI subtraction algorithm showing tolerance to saturated pixels and lack of ringing at nulls of the reference PSF.

REFERENCES

- [1] Marois, C., Doyon, R., Nadeau, D., Racine, R., and Walker, G. A. H., “Effects of Quasi-Static Aberrations in Faint Companion Searches,” **8**, 233–243 (2003).
- [2] Baudoz, P., Boccaletti, A., Baudrand, J., and Rouan, D., “The Self-Coherent Camera: A new tool for planet detection,” *Proceedings of the International Astronomical Union* **1**, 553–558 (Oct. 2005).
- [3] Guyon, O., “Imaging Faint Sources within a Speckle Halo with Synchronous Interferometric Speckle Subtraction,” **615**, 562–572 (Nov. 2004).
- [4] Baudoz, P., Boccaletti, A., Baudrand, J., and Rouan, D., “The Self-Coherent Camera: a new tool for planet detection,” in [*IAU Colloq. 200: Direct Imaging of Exoplanets: Science & Techniques*], Aime, C. and Vakili, F., eds., 553–558 (Jan. 2006).
- [5] Bordé, P. J. and Traub, W. A., “High-Contrast Imaging from Space: Speckle Nulling in a Low-Aberration Regime,” **638**, 488–498 (Feb. 2006).
- [6] Give’On, A., Belikov, R., Shaklan, S., and Kasdin, J., “Closed loop, DM diversity-based, wavefront correction algorithm for high contrast imaging systems,” *Optics Express* **15**, 12338 (Jan. 2007).
- [7] Serabyn, E., Wallace, J. K., and Mawet, D., “Speckle-phase measurement in a tandem-vortex coronagraph,” **50**, 5453 (Oct. 2011).
- [8] Sauvage, J. F., Mugnier, L., Paul, B., and Villicroze, R., “Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging,” *Optics Letters* **37**, 4808 (Dec. 2012).
- [9] Gerard, B. L., Marois, C., and Galicher, R., “Fast Coherent Differential Imaging on Ground-based Telescopes Using the Self-coherent Camera,” *AJ* **156**, 106 (Aug. 2018).
- [10] Potier and et al., “Improving VLT/SPHERE without additional hardware: First on-sky calibration of the quasi-static aberrations with the dark hole technique,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *current proceedings*, **12185-236** (2022).
- [11] Gerard, B. L., Marois, C., Galicher, R., and Véran, J.-P., “Fast focal plane wavefront sensing on ground-based telescopes,” **0703**, 1070351 (July 2018).
- [12] Lardière, O. and et al., “Optical design of SPIDERS, a Subaru Pathfinder Instrument for Detecting Exoplanets & Retrieving Spectra,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *current proceedings*, **12185-156** (2022).
- [13] Marois, C. and et al., “Deployment of focal plane WFS technologies on 8-m telescopes: from the Subaru SPIDERS pathfinder, to the facility-class GPI 2.0 CAL2 system,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *current proceedings*, **12185-70** (2022).

- [14] Marois, C., Gerard, B., Lardière, O., Anthony, A., Bradley, C., Dunn, J., Fu, Q., Hardy, T., Heidrich, W., Herriot, G., Nielsen, E., Sivanandam, S., Sivransky, D., Thibault, S., Thompson, W., and Véran, J.-P., “Upgrading the Gemini Planet Imager calibration unit with a photon counting focal plane wavefront sensor,” in [*Adaptive Optics Systems VII*], **11448**, 1144873, International Society for Optics and Photonics (Dec. 2020).
- [15] Lardière, O., Gerard, B., Thompson, W., Marois, C., Véran, J.-P., Blain, C., Heidrich, W., and Fu, Q., “Optical design and preliminary results of NEW EARTH, first Canadian high-contrast imaging laboratory test bench,” in [*Adaptive Optics Systems VII*], **11448**, 114486Y, International Society for Optics and Photonics (Dec. 2020).
- [16] Singh, G., Martinache, F., Baudoz, P., Guyon, O., Matsuo, T., Jovanovic, N., and Clergeon, C., “Lyot-based Low Order Wavefront Sensor for Phase-mask Coronagraphs: Principle, Simulations and Laboratory Experiments,” **126**, 586 (June 2014).
- [17] Singh, G., Lozi, J., Guyon, O., Baudoz, P., Jovanovic, N., Martinache, F., Kudo, T., Serabyn, E., and Kuhn, J., “On-Sky Demonstration of Low-Order Wavefront Sensing and Control with Focal Plane Phase Mask Coronagraphs,” **127**, 857 (Sept. 2015).
- [18] Singh, G. and et al., “Pupil-plane LLOWFS simulation and laboratory results from NEW-EARTH’s high-contrast imaging testbed,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *current proceedings*, **12185-192** (2022).
- [19] Johnson, A. B. and et al., “Blinking the fringes, initial development and results of the Ultra-Low Speed Optical Chopper for the Self-Coherent Camera,” in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], *current proceedings*, **12185-271** (2022).
- [20] Gerard, B. L., *Exoplanet Imaging Speckle Subtraction: Current Limitations and a Path Forward*, thesis (2020).
- [21] Wahhaj, Z., Z., Milli, J., Romero, C., Cieza, L., Zurlo, A., Vigan, A., Peña, E., Valdes, G., Cantalloube, F., Girard, J., and Pantoja, B., “A search for a fifth planet around HR 8799 using the star-hopping RDI technique at VLT/SPHERE,” *A&A* **648**, A26 (Apr. 2021).
- [22] Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., and Nadeau, D., “Angular Differential Imaging: A Powerful High-Contrast Imaging Technique*,” *ApJ* **641**, 556 (Apr. 2006).
- [23] Walker, G., Chapman, S., Mandushev, G., Racine, R., Nadeau, D., Doyon, R., and Véran, J. P., “Shades of Black: Searching for Brown Dwarfs and Giant Planets,” **56**, 449 (Jan. 1999).
- [24] Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., and Marois, C., “Speckle Noise and the Detection of Faint Companions,” *PASP* **111**, 587 (May 1999).
- [25] Marois, C., Doyon, R., Racine, R., and Nadeau, D., “Efficient Speckle Noise Attenuation in Faint Companion Imaging,” *PUBL ASTRON SOC PAC* **112**, 91–96 (Jan. 2000).